


Securing AI Agents by Design




Overview

Agentic AI represents the next stage in enterprise AI, advancing beyond simple chatbots to intelligent software agents capable of perceiving, reasoning, planning multi-step goals and executing actions through external tools or APIs. Major cloud platforms like Google AgentSpace, AWS Agents and Microsoft AgentFlow have introduced tailored solutions, while self-hosted frameworks such as LangChain, LangGraph and Semantic Kernel are gaining traction among enterprises. However, the autonomy of these agents introduces new risks by expanding the attack surface. IDC forecasts that over 40% of enterprise workloads will utilize autonomous agents by 2027 (up from under 5% in 2024). Businesses must prioritize security and governance to navigate this rapid evolution.

Challenges

The emerging threat landscape

The rise of AI agents and Large Language Models (LLMs) has unveiled a wave of cybersecurity risks. The OWASP Top 10 for LLM and Agentic AI highlights vulnerabilities such as prompt injection, excessive agency, supply-chain compromise and model stealing, which in the context of Agentic AI translate to tangible business risks, including:



Data leaks: A rogue “exfiltration” tool-call could expose proprietary trading strategies.



Compromised reasoning loops: A poisoned dependency could hijack an agent’s workflow.



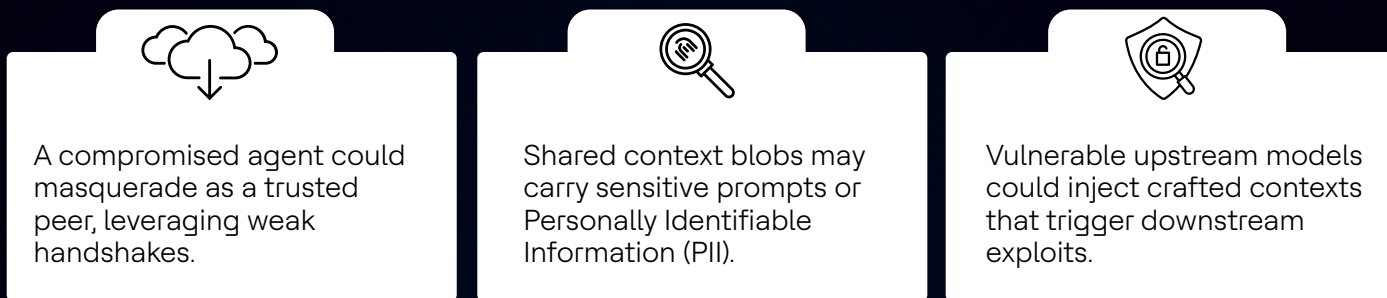
Fraudulent operations: Over-permissioned customer-service bots might execute unintended actions, like filing counterfeit refunds.

These risks aren’t theoretical. They represent the new threat landscape of autonomous AI agents, demanding proactive mitigation strategies to ensure secure, scalable deployment.



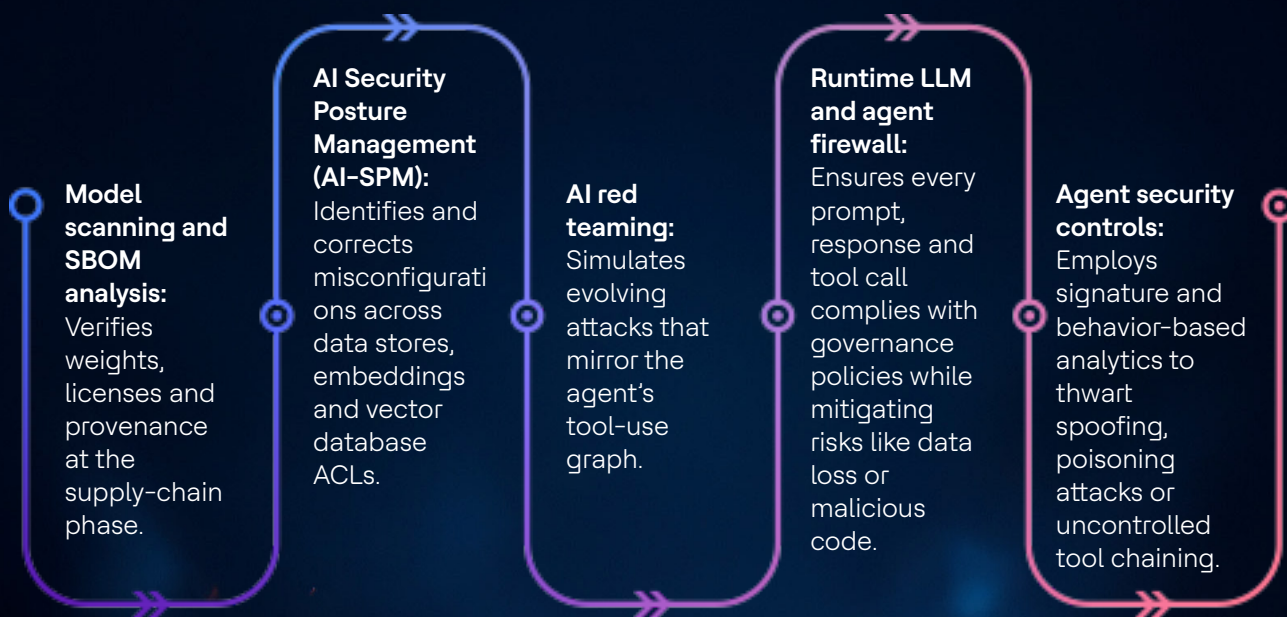
A2A and MCP - why new protocols need new guard-rails

Emerging protocols such as Agent-to-Agent (A2A) and Model-Context Protocol (MCP) aim to standardize how agents discover, negotiate and invoke each other's capabilities. While these protocols promise seamless Lego-block-like composability, they also amplify security risks, such as:



Our Solution

HCLTech, together with Google Cloud and Palo Alto Networks, addresses these gaps. Key capabilities include:



HCLTech provides security consulting, turnkey reference architectures and Universal Managed Detection and Response (UMDR) services enabled by these solutions to deliver secure Agentic AI behavior. This allows businesses to focus on specific goals and outcomes, not complexity.



Use-cases: Secure AI-by-design in action

Most organizations move through four critical stages when adopting Agentic AI solutions.

- 1 Experimentation:** Sandboxed testing of agents and frameworks. Palo Alto Networks Prisma AIRS flags early-stage vulnerabilities (e.g., licensing flaws or unsafe code behaviors) while operating in “observe-only” mode during the evaluation phase.
- 2 Pilot deployment:** Launching small-scale use cases like customer-support chatbots or internal FAQ agents. HCLTech configures least-privilege tool access, sets up SBOM attestation gates and maps policies to reduce exposure.
- 3 Production rollout:** Scaling agents into revenue-critical workflows. Runtime enforcement activates here, blocking threats like unauthorized tool chaining, excess API calls and entity spoofing.
- 4 Operational monitoring and optimization:** Continuous monitoring of security threats, policy drifts, operational inefficiencies or cost overruns. Palo Alto Networks Prisma AIRS seamlessly integrates with HCLTech’s XDR/SIEM telemetry for proactive optimization and long-term governance.

Building a secure agentic AI future

Agentic AI is poised to transform critical fields, viz., orchestrating global supply chains, managing hedge portfolios, enabling automated triage in healthcare and much more. However, security is non-negotiable. By embedding security and governance by design into the foundation of AI systems and maintaining it throughout development, deployment, scaling and optimization, HCLTech and Palo Alto Networks empower organizations to confidently adopt autonomous AI agents, ensuring that they operate, evolve to learn and deliver value.

Get started today

Ready to explore Secure Agentic AI for your next initiative? Reach out to your HCLTech or Palo Alto Networks representative to schedule a workshop.



HCLTech | Supercharging
Progress™

hcltech.com